DARPA-BAA-14-60
STAC
Frequently Asked Questions

As of September 30, 2014

Q75: During the engagements, will the TA1 and TA3 teams be required to also analyze obfuscated bytecode? Can teams assume that the code was generated from a standard Java compiler (e.g., javac)?

A75: The BAA encourages TA2 AC proposers to consider what kinds of benign but difficult-to-analyze kinds of code might hinder the efforts of an automated analysis to determine what a program does and thus make it difficult to identify algorithmic resource usage vulnerabilities in that program. It further encourages TA2 AC technical approaches that seek to produce a collection of challenge programs that cover differing levels of difficulty. The use of obfuscation and bytecode generated by means other than a standard Java compiler is in scope for STAC. TA2 AC proposers may choose to incorporate these techniques into their technical approach or not, as they see fit.

Q74: Why are there additional TA2 AC Software Deliveries on months 18 and 36? These don't seem to fit the regular engagement cycle. Is this a typo in the BAA?

A74: The BAA is correct, but may deserve some explanation: Month 18 is the end of Phase 1. Month 36 is the end of Phase 2. The Government needs a delivery of the software created by the TA1 R&D performers and TA2 AC performers at the end of each phase for programmatic reasons. For example, DARPA may choose not to exercise the Phase 2 options of some TA1 R&D and TA2 AC performers, and will need a delivery of their software as it exists at the end of Phase 1. For TA1 R&D performers, the ends of Phase 1 and Phase 2 correspond to the months they would be delivering software in the regular engagement cycle anyways. For TA2 AC performers, the ends of Phase 1 and Phase 2 do not line up with the regular engagement cycle, and consequently they have two additional software deliveries due at the ends of these phases. DARPA expects that these two additional TA2 AC software deliveries on months 18 and 36 will be a snapshot of the state of the software's development in progress at that time, as the TA2 AC performers should be working towards their regular month 20 and 38 deliveries as the phases are ending.

Q73: Can faculty from other countries be involved with a team?

A73: Yes.

Q72: Are subcontractors allowed to team with multiple primes for the same TA? If so, are there any stipulations?

A72: The BAA states that "While proposers may submit proposals for all Technical Areas, they should not expect DARPA to select more than one of these proposals." This warning applies to both prime contractors and subcontractors - if a prime or subcontractor appears on proposals to two or more TAs, they should not expect DARPA to select more than one of those proposals. This warning does not apply to proposals within the same TA. DARPA will consider selecting multiple proposals within the same TA that share a subcontractor. In this situation, DARPA expects potential subcontractors to offer different teams of key personnel in each proposal.

Q71: If we submit a proposal to grants.gov, do we also need to register in DARPA BAA site?

A71: No. If you submit via grants.gov, you do not need to worry about the DARPA BAA site.

Q70: Page 2 of the BAA states the program will be “42” months, is that a typo?

A70: Yes, the program will run 48 months! Thanks for pointing that out. The BAA has been amended to correct page 2.

Q69: Will there be a DD254 required to be attached to the BAA?

A69: No.

Q68: TA1 and TA2 intrinsically have a conflict of interest with each other. Is there any conflict of interest mitigation strategy that would realistically be acceptable to DARPA, PM, Etc.? For example, on R&D group at company X is interested in TA1 and a different R&D group at X is interested in TA2, and they firewall themselves from each other in some way demonstrable to DARPA?

A68: No.

Q67: Are you going to provide a Proposer’s Day attendee list?

A67: No.

Q66: Does the posted budget include or exclude agent fees, SETA costs, Studies, etc.?

A66: The posted budget represents the funds available to performers.

Q65: Is human use expected in TA3?

A65: No, Human use is NOT expected.

Q64: Given that TA1 analyses will be conducted on Java bytecode. Will speed metrics be measured with respect to the KLOC of original Java source that produced the bytecode or something like KLO-bytecode?

A64: DARPA will determine precisely what method will be used to measure speed after the program has begun. However, DARPA anticipates that the metrics will be generally as described in the BAA, and that finding analysis techniques and tools that can enable an analyst to analyze larger programs faster will remain a program goal.

Q63: Will TA1 performers get source for Java bytecode AFTER competitive engagements:

A63: DARPA will determine this after the program has begun.

Q62: Is the PM anticipating the fostering of primes and subs pairings (based on proposal contents)?

A62: No.

Q61: Is the program transformation of interest? E.g. rather than just stating a program has a timing channel, would it be useful/preferred to suggest a transformation that would eliminate that channel?

A61: No. Approaches that seek to eliminate vulnerabilities without first determining whether or not they are present by simply transforming a program into a form in which vulnerabilities cannot occur are out of scope for STAC.

Q60: What machine model is assumed? Different instructions might take different amounts of time on different VMs. The JIT might act differently on different VMs. Should we assume a particular VM semantics? If not, what semantics should we assume?

A60: DARPA will make this determination after the start of the program.

Q59: Do the solutions to the two problems – complexity attacks and side channels need to be related? Is it preferred if they are?

A59: Your proposal must address both problems, but your solutions to the two do not need to be related.

Q58: What side channels are of interest?

A58: Only side channels based on the space and time resource usage of algorithms are in scope. Side channels based on physical characteristics such as power consumption, thermal emissions or the emissions of other kinds of radiation are out of scope. Side channels based on differential error behavior are out of scope.

Q57: Will the competitions be on a single site?

A57: As described in the BAA, engagements may be made up of several parts. Some parts will be like take-home exams. Some parts will be held live at PI meetings. The take-home exams will occur at separate sites – specifically, at each performer's site. The live parts will be at a single site – specifically, the site of the PI meeting.

Q56: Will Engagements be at the same time or adjacent?

A56: DARPA anticipates all TA2 R&D and TA3 Control performers participating in the engagements simultaneously.

Q55: How long will Engagements last?

A55: As described in the BAA, engagements may be made up of several parts. Some parts will be like take-home exams. Some parts will be held live at PI meetings. DARPA anticipates allowing performers some number of days or weeks to complete the take-home exams. DARPA anticipates the live parts will last some number of hours.

Q54: The BAA stipulates Java bytecode. Is there any guidance on underlying processing architecture (hardware)?

A54: No guidance.

Q53: What does the schedule look like for TA3 performers? Are they only there for challenge events or can they spend time in between preparing for challenges?

A53: The TA3 Role is described in the BAA. You may propose activities and a level of effort you believe is appropriate and fitting.

Q52: How is user/sensitive secret input identified? Is the onus of identifying such inputs on the tool provider?

A52: The onus is on the tool provider.

Q51: Is the criterion for a Denial of Service (DOS) left to the tool implementer or are there hard timing constraints (microseconds/seconds)?

A51: DARPA expects to provide guidance on this issue after the program begins. DARPA anticipates this guidance to be symbolic and in terms of adversary vs. defender workload. For example, adversary work no worse than O(n) and defender work no better than O(n2).

Q50: Are detecting DOS vulnerabilities related to dead-lock/live-lock, race conditions etc… in scope as well?

A50: No.

Q49: BAA states SOW by year; should this be by phase?

A49: Yes, this should be by phase. The BAA has been amended to correct this issue.

Q48: Of the three types of security properties (confidentiality, integrity, and availability), the BAA deals with confidentiality (side-channels) and availability (resource usage attacks). Do you see any challenges in terms of (data) integrity?

A48: Only algorithmic complexity attacks (availability) and side channels (confidentiality) are in scope for STAC.

Q47: TA2 – What is the stance on code reuse for generating challenge bytecode?

A47: Please see the discussion of the “looking for differences” shortcut in the description of TA2 in the BAA. You are encouraged to propose a viable solution.

Q46: For programs with potential side-channel vulnerabilities, will you specify what part of program-input constitutes a secret that should be protected?

A46: No.

Q45: Can TA1 proposal address only algorithmic complexity related DOS attack vulnerabilities and not address side-channel attacks and vice-versa?

A45: No. Both must be addressed.

Q44: A program with linear space/time complexity may also be subject to DOS attack, if subjected to extremely large input, e.g. by supplying a terabyte long file to a file search algorithm. Should we worry about such scenarios as well?

A44: No. We're looking for cases where the defender's workload exceeds the adversary's workload, for example O(n) vs. O(n2) space or time resource usage.

Q43: Are you open to funding the development of practical and new static analysis techniques and tools which contradict some of the theoretical feasibility assumptions in the BAA?

A43: DARPA welcomes proposals for innovative approaches that enable revolutionary advances.

Q42: The second interest area relates to side-channel attacks. Is the interest in this area primarily remotely accessible approaches like timing or does it include local, approaches, like power analysis also?

A42: Only side channels based on the space and time resource usage of algorithms are in scope. Side channels based on physical characteristics such as power consumption, thermal emissions or the emissions of other kinds of radiation are out of scope. Side channels based on differential error behavior are out of scope. DARPA encourages proposers to focus on software that will be deployed in scenarios where an adversary has the ability to directly or indirectly provide inputs to that software and to directly or indirectly observe that software's outputs in order to exploit algorithmic resource usage vulnerabilities.

Q41: Are there any specifications or constraints on the analysis platform?

A41: As described in the BAA, DARPA is seeking semi-automated approaches capable of detecting new vulnerabilities.

Q40: What is the size/complexity of the code base we will be analyzing?

A40: DARPA will make this determination after program start after considering the approaches proposed by the TA2 AC and TA4 EL performers.

Q39: What are the target libraries and programs (cryptographic libraries, file/expression, parsers, etc)?

A39: DARPA will make this determination after program start after considering the approaches proposed by the TA2 AC and TA4 EL performers.

Q38: Will there be native libraries/code called by the target application?

A38: DARPA discourages TA2 AC approaches that rely on the use of native code. Only bytecode for the Java Virtual Machine is in scope for analysis in STAC.

Q37: What is the reference JVM that will be used for running applications?

A37: DARPA will make this determination after program start.

Q36: Most of the vulnerabilities described in the literature are found in systems or libraries developed in languages other than Java. Do we only target systems developed in Java?

A36: Only bytecode for the Java Virtual Machine is in scope for analysis in STAC.

Q35: There are two main performance axes: scale and speed. Do we measure the proposed techniques against current (state of art) static analysis techniques ability to handle larger programs?

A35: DARPA anticipates that engagements will be a kind of "horse race" between all TA1 R&D performers and the TA3 Control performer that will measure the performance of these performers relative to each other. The TA3 Control performer will represent the current state of art.

Q34: TA1/TA4: If performance is measured with respect to KLO-bytecode, will performers be responsible for somehow providing a solution for translating these numbers into metrics comparable to the KLOC in the cited works?

A34: DARPA will make this determination after program start after considering the approaches proposed by the TA4 EL performer.

Q33: Can we propose to TA3 and TA1? What is the expected schedule to look like for TA3 (what do they do between engagements).

A33: Offerors may respond to multiple technical areas by submitting multiple proposals, one for each technical area. Such offerors should not expect DARPA to select and fund more than one of these proposals. The Government will make the choice of which, if any, proposal to select and fund. The TA3 Role is described in the BAA. You may propose activities and a level of effort you believe is appropriate and fitting.

Q32: Is DARPA interested in detecting programs that slowly leak resources (due to mismanagement) without the program being given any special inputs?

A32: The detection of memory leaks is not in scope for STAC. Please see the public descriptions of the CRIME and BREACH attacks cited in the BAA for examples of side channels.

Q31: Will legitimate adversary observation points for side channels be clearly circumscribed in advance?

A31: DARPA will make this determination after program start after considering the approaches proposed by the TA4 EL performer.

Q30: What is the minimal bandwidth for side-channels that should be considered? (BPS)

A30: DARPA expects to provide guidance on this issue after the program begins.

Q29: Are multi-threaded programs and the associated locking mechanisms in scope?

A29: Programs that meet the STAC scope restrictions are in scope regardless of whether or not they are multi-threaded. Searching for deadlock, livelock, and race conditions is out of scope for STAC.

Q28: We have to model the time/resource characteristics of external modules/APIs. Will algorithmic flaws in these external components be excluded from the metrics and challenges?

A28: Any component expressed as bytecode for the Java Virtual Machine is in scope for analysis in STAC. DARPA anticipates making it clear which components are and are not in bounds for analysis ahead of each engagement.

Q27: To what extent should system factors be accounted for (e.g. JIT complication)? Or system loading – other processes?

A27: Because TA1 R&D and TA3 Control performers will be analyzing real programs, it seems reasonable to expect that system factors such as load may impact some analysis approaches. The extent of this impact seems likely to depend on the details of the specific approach.

Q26: How broad are the definitions of time and space? For example is space more than message size (e.g. memory consumed). Likewise time could be algorithmic on a CPU or wait time on a complex database query.

A26: The definitions are broad. Space may be message size or memory consumed depending on the details of the program being analyzed.

Q25: Are side channels limited to space and time or can others signals such as power usage be used?

A25: Only side channels based on the space and time resource usage of algorithms are in scope. Side channels based on physical characteristics such as power consumption, thermal emissions or the emissions of other kinds of radiation are out of scope. Side channels based on differential error behavior are out of scope.

Q24: Across a 4-year period, the "present day" methods available to TA3 will: 1) Track with published methods. 2) Remain static, 3) Other.

Q24: It can track.

Q23: Should one analysis program find in both DOS problems and Side-channel problems?

A23: The number of programs tools is up to you. But your project must address both problems.

Q22: Are you looking for analysis beyond software (e.g. networks, etc..)

A22: The focus is on programs. Programs can use networks, but we'll be looking for problems in code rather than evidence of problems in net traces. You can look at traces, but we want to know about the vulnerabilities in the code.

Q21: Side channel vulnerabilities may arise due to information inadvertently placed on observable "signal" Time and space usage are only two examples of signals. Will STAC consider other signals such as power consumption, physical outputs, or the actual content of error messages?

A21: Only side channels based on the space and time resource usage of algorithms are in scope. Side channels based on physical characteristics such as power consumption, thermal emissions or the emissions of other kinds of radiation are out of scope. Side channels based on differential error behavior are out of scope.

Q20: Is STAC open to hardware-assisted techniques?

A20: Yes. You might conceivably program an FPGA to do computations related to automated analysis, for example.

Q19: How do you foresee the tests playing out?

A19: As described in the BAA, engagements may be made up of several parts. Some parts will be like take-home exams. Some parts will be held live at PI meetings. The take-home exams will occur at separate sites – specifically, at each performer's site. The live parts will be at a single site – specifically, the site of the PI meeting. DARPA anticipates allowing performers some number of days or weeks to complete the take-home exams. DARPA anticipates the live parts will last some number of hours.

Q18: Do you have the code base we will be analyzing?

A18: No. It will be produced by the TA2 AC performers.

Q17: When will we get the code base to analyze?

A17: Please see the delivery schedule described in the BAA.

Q16: What do you anticipate that an acceptable false negative/false positive rate will be (other than as good as possible)?

A16: As good as possible.

Q15: The STAC webpage on federalgrants.com website lists the expected number of awards to be 20. Given that TA 2-4 may receive a total of 6 awards, can we assume that up to 14 awards may be made for TA1?

A15: The above website is not associated with the Government. DARPA will not speculate on how they came up with their estimated number of awards.

Q14: Are you planning on a program with a small number of teams (that may be large) or a large number of (possibly small) teams?

A14: DARPA anticipates making this determination based on the proposals received.

Q13: Are there any publication restrictions?

A13: STAC is a 6.1 fundamental research program. As stated in the BAA, DARPA does not anticipate publications restrictions. Please see the BAA for details.

Q12: We plan to propose in TA1. If the UARC affiliated with our university is proposing in TA2, will there be a COI?

A12: Please submit these kinds of COI-related questions to STAC@darpa.mil and include a summary of the potential conflict. Specific COI questions will be addressed directly with the submitter vs. the FAQ process.

Q11: The term "Java bytecode" used throughout the BAA is somewhat ambiguous because the same Java source code can be compiled into various bytecode formats (perhaps the two most notable examples, the Android JVM and Oracle's hotspot JVM, have vastly different bytecode formats). What is the target bytecode format of interest to STAC?

A11: The STAC program is concerned with resource usage vulnerabilities in programs expressed in bytecode for the Java Virtual Machine. Vulnerabilities in programs expressed in languages other than this are out of scope. For example, vulnerabilities in programs expressed in Dalvik bytecode are out of scope.

Q10: Can you narrow down the domain for the types of software STAC will be expected to analyze (mobile applications, server components, desktop applications)?

A10: The STAC program is concerned with resource usage vulnerabilities in programs expressed in bytecode for the Java Virtual Machine. In addition, deployment scenario is a more important criterion than type. DARPA encourages proposers to focus software that will be deployed in scenarios where an adversary has the ability to directly or indirectly provide inputs to that software and to directly or indirectly observe that software's outputs in order to exploit algorithmic resource usage vulnerabilities. Typical mobile, server, and desktop applications may all be deployed in these scenarios, particularly if they are designed to communicate over the Internet.

Q9: Is it reasonable for a proposer to assume that additional non-source testing artifacts such as compiled automated test code (e.g., Junit test suites) and/or example inputs will be included with the challenge problems?

A9: No.

Q8: Aren't side channels dependent on the hardware where the Java software will execute (e.g. smart cards)?

A8: Only side channels based on space and time resource usage are in scope for STAC. These should be hardware-independent.

Q7: What outputs will a TA1 team provide to indicate that they have identified a side channel? Will they identify the information that is leaked; the bandwidth of the leak; the noise, if any, in the channel; and the portion of the software responsible for the leak? What outputs will a TA1 team provide to indicate that they have identified an algorithmic complexity vulnerability? Input data that causes a challenge problem to run slowly or crash may or may not be a symptom of a complexity attack.

A7: STAC is motivated by the need to find and remediate algorithmic resource usage vulnerabilities. To serve that need, we expect both the existence of the vulnerability and the portion of the software responsible to be of interest. We expect to work out the precise details of how performers will report their results in the engagements once the program begins. You are welcome to propose a method.

Q6: Would proposals from two completely separate Limited Liability Corporations (LLCs) owned by a single large company be considered as coming from one proposer in the context of Section III.D.1 of the BAA?

A6: Yes.

Q5: Is differential power analysis in-scope as a tool for side channel creation?

A5: No.

Q4: Is it in scope for challenge problems to include hidden content that is obviously deliberately malicious once discovered?

A4: A vulnerability is a vulnerability regardless of the intent with which it was put in the software.

Q3: Physical side channel attacks. Side channel attacks invariably exploit information that leaks through the computation platform in the form of physical quantities: energy, temperature, current, radiation, etc. Are these type of attacks within scope?

A3: Only side channels based on space and time resource usage are in scope. Side channels based on energy, temperature, current, and radiation are out of scope.

Q2: Software for cyber-physical systems manipulates information that originates in the physical world and that is fed back to the physical world via actuators. Does this type of software fit within STAC's scope?

A2: DARPA encourages proposers to focus on software that will be deployed in scenarios where an adversary has the ability to directly or indirectly provide inputs to that software and to directly or indirectly observe that software's outputs in order to exploit

algorithmic resource usage vulnerabilities. It is possible to imagine some cyber-physical systems that might fall into this category, perhaps some of it might be encoded as bytecode for Java VMs. Note that in each engagement, all TA1 R&D performers will use their techniques and tools to search for algorithmic resource usage behavior vulnerabilities in a set of challenge programs chosen by the TA4 EL performer from those produced by the TA2 AC performers. Thus the TA4 EL and TA2 AC performers - not the TA1 R&D performers - will determine what kind of challenge programs the TA1 R&D performers will analyze in engagements.

Q1: May a proposer receive an award as a Prime in one Technical area, and an award as a subcontractor on a team in another Technical Area?

A1: The BAA states that "While proposers may submit proposals for all Technical Areas, they should not expect DARPA to select more than one of these proposals." This warning applies to both prime contractors and subcontractors - if a prime or subcontractor appears on proposals to two or more TAs, they should not expect DARPA to select more than one of those proposals.